A Bayesian Treatment of Risk for Radiation Hardness Assurance[1]

R. Ladbury *Member, IEEE*, J. L. Gorelick, *Member, IEEE*, M. A. Xapsos, *Member, IEEE*,
T. O'Connor, and S. Demosthenes, *Member, IEEE*

***Abstract*— We construct a Bayesian risk metric with a method that allows for efficient and systematic use of all relevant information and provides a rational basis for RHA decisions in terms of costs and mission requirements.**

***Index Terms*—radiation effects, reliability estimation, statistics, quality assurance, probability.**

I. INTRODUCTION

Because economic and schedule constraints often preclude radiation testing with samples large enough for rigorous statistical inference, Radiation Hardness Assurance (RHA) decisions are often based on a combination of test data with simulation results, technical information and expert opinion. Efficient combination of all of these types of information is particularly important for critical components with marginal radiation hardness. In reference 1 we discussed use of archival radiation data to supplement data from radiation lot acceptance testing (RLAT) and provide a sufficient statistical basis for inference.[1] Here we examine whether RHA decisions may be amenable to a Bayesian statistical treatment—often the most efficient treatment of diverse types of information when decisions must be made with limited data.

In the Bayesian paradigm, one uses archival data, expert opinion and all other information available prior to testing to construct a Prior probability distribution (or "Prior") $P(\boldsymbol{A})$ for a hypothesis ($\boldsymbol{A}$), summarizing our *a priori* expectations of component performance. We then update this distribution with test data $\boldsymbol{B}$ using Bayes' Theorem:

$$P(A\,|\,B) = \frac{P(B|A)\times P(A)}{(P(B|A)+P(B|{\sim}A))} \qquad (1),$$

where $P(\boldsymbol{A}|\boldsymbol{B})$ (read probability of **A** given $\boldsymbol{B}$) is our updated distribution given $\boldsymbol{B}$, $P(\boldsymbol{B}|\boldsymbol{A})$ is the likelihood of observing $\boldsymbol{B}$ given our Prior P(A) and the denominator is a normalizing factor representing the likelihood of observing B whether our original hypothesis (A) was correct or not (~A—read "not A"). (Reference 2 is a good introduction to Bayesian reasoning.)

In Bayesian statistics, provided the test data B are consistent with hypothesis A, the information in P(A) supplements B, yielding greater confidence in inferred performance.

Elements of a Bayesian treatment have been suggested previously.[3] Moreover, a Bayesian treatment is natural, since many aspects of hardness assurance resemble Bruno de Finetti's classic problem on Bayesian probability. (See figure 1.) In this problem, several urns contain different proportions of white and black stones, representing success and failure, respectively.[4] One selects an urn, and only then is the probability distribution defined for the problem. That is, the probability distribution itself is a random variable. This example is analogous to RHA situations in which a randomly selected flight-lot determines mission success probabilities. Similarly, for short-duration missions, the mission radiation environment—which determines both degradation and single-event behavior—can also be viewed as a random variable, defined only when one selects a launch date.

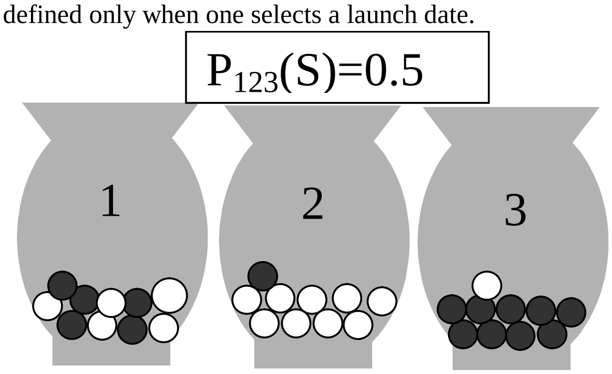

Fig. 1 In Bruno de Finetti's thought experiment, the probability of drawing a black or white stone is not defined until one first selects an urn (or flight lot or launch date). The probability distribution itself is a random variable.

II. BAYESIAN FORMALISM AND MINIMIZING RISK

In a typical Bayesian problem we have N independent applications of a given part type. The parts are vulnerable to a

Manuscript received September 16, 2005. This work was supported in part the NASA James Webb Space Telescope (JWST) Program and the NEPP Electronic Radiation Characterization (ERC) Project.

R. Ladbury is with NASA/GSFC, Greenbelt, MD 20771, USA (phone: 301-286-1030; fax: 301-286-4699; e-mail: Raymond.L.Ladbury.1@gsfc.nasa.gov).

J. L. Gorelick is with Boeing Space and Intelligent Systems (BSIS), Los Angeles, CA USA, (jerry.l.gorelick@boeing.com).

M. A. Xapsos is with NASA/GSFC, Greenbelt, MD 20771, USA (e-mail: Michael.A.Xapsos.1@gsfc.nasa.gov).

T. O'Connor is with Ball Aerospace and Technologies Corp., Boulder, CO, 8 , USA (email: toconnor@ball.com)

S. Demosthenes is with Ball Aerospace and Technologies Corp., Boulder, CO, 8 , USA (email: sdemosth@ball.com)

[1] This work was supported by the NASA James Webb Space Telescope (JWST) Program and the NEPP Electronic Radiation Characterization (ERC) Project

failure mechanism induced by some stress x (e.g. TID, LET, etc.), and all parts obey the same failure distribution $p_f(x)$. We must decide whether to use the parts "as is" or to require remediation to increase the probability that the part meets its requirements.[5,6] Two risks are inherent in this decision:

1) We may decide to fly the part "as is" when its failure probability is unacceptably high; or

2) We may decide the part requires remediation when it would have succeeded "as is".

Defining risk as the product of an outcome's probability and its cost, Risk 1—the failure risk, R_f — is the product of the failure probability P_f and the failure cost C_f. Recognizing that C_f may be time dependent (e.g. an early failure may be more costly than a late one), we can express R_f in terms of the failure probability density function (pdf) p_f as

$$R_f = \int_0^{X_m} p_f(x) \times \sum_{i=1}^{N} C_{f\,i}(x) dx \qquad (2)$$

In (2), the sum is over the N applications of the part, and p_f and C_f are functions either of time (x=t, X_m= T_m—the mission duration) or the stress (e.g. TID, so that x=dose and X_m = mission dose). Since typically C_f is a function of time and p_f is a function of the stress, this implies that we have a way of treating the time variability of the stress—at least in a "worst-case" way or for a given confidence level. We discuss such a treatment below. On the other hand, if C_f is constant, R_f = C_fP_f, where P_f is the cumulative mission failure probability.

Similarly, Risk2—the over-remediation risk R_r—is realized if the part would have succeeded "as is" and we implement remediation at cost C_r. Then $R_r = P_sC_r = (1- P_f)C_r$. We adopt the strategy that minimizes risk—continuing remediation until $R_f < R_r$. (See figure 2.) To better define the failure probability, we may conduct a test at a cost C_t. However, since a decision to test automatically incurs a cost C_t, we should only test if the resulting risk reduction is likely to exceed C_t.

Effective tests must address the main sources of uncertainty in the failure probability. For total ionizing dose (TID) and displacement damage (DD), the main uncertainties are usually lot-to-lot and part-to-part variability, so TID and DD tests usually involve measuring the response of more parts to tighten bounds on variability. Probabilities for single-event effects (SEE) scale with SEE rates, and SEE rate uncertainties are usually dominated by systematic errors in the SEE rate calculation methods and Poisson fluctuations in the observed error counts that determine SEE cross sections. Since the systematic errors are inherent, SEE rate/probability uncertainties are often best reduced by increasing the counts on which SEE cross sections are based. This is especially true for destructive SEE and single-event functional interrupts (SEFI), where the disruptive and costly nature of the error makes it difficult to accumulate large error counts.

III. Defining Costs

The costs of testing, remediation and failure (C_t, C_r , C_f, respectively) determine where remediation and testing may end. Thus, successful application of the method depends on developing realistic estimates for C_f.

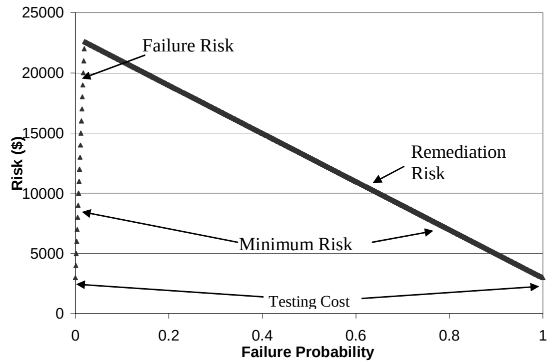

Fig. 2 Taking risk as the minimum arising due to failure or unnecessary remediation defines a cost scale for the RHA effort.

Sometimes, contractual penalties or incentives impose a cost structure. Otherwise, C_f will be based on the functions compromised by the failure and by how long the effects of the error persist. Determining a cost for all the potential error or failure mode can be simplified by binning the errors/failures into severity categories (from trivial to mission critical) and basing cost on severity. In this case, the analysis resembles procedures often followed for SEE or reliability analyses (such as a FMECA, or Failure Modes, Effects and Criticality analysis), except that one assigns a specific cost to each error/failure category.[7] (See figure 3.) Failures and degraded performance are usually permanent, although in redundant systems, a failure's consequences may last only for the time needed to replace the failed unit with a redundant one.

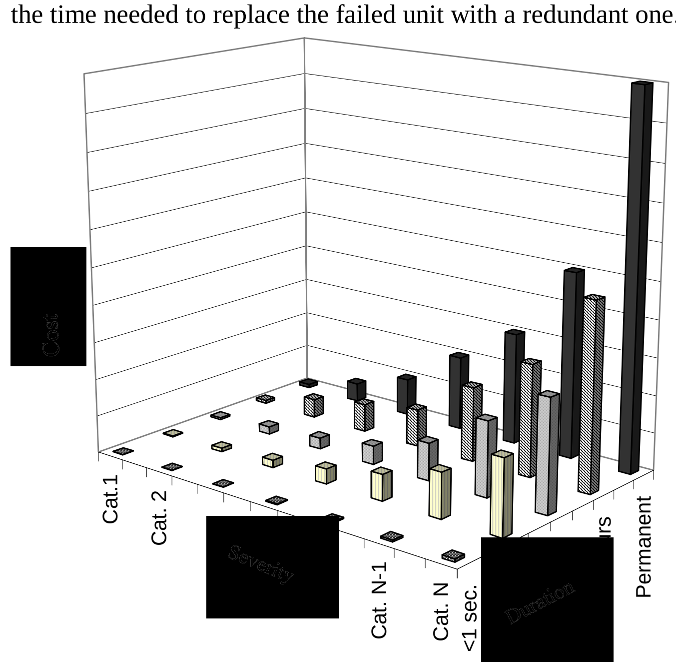

Fig. 2 Costs of a radiation-induced error or failures depend on what functions are lost and for how long. Failures due to TID or displacement damage are usually permanent.

Application, mission or contract requirements also determine the time dependence of C_f. Some failures have the same consequences no matter when they occur. In other cases, C_f may decrease with time (more or less linearly) as mission objectives are completed. For complicated time dependence, equation (2) may require numerical integration.

Finally, it should be noted that failures may have additional costs (e.g. insurance premiums, reputation costs, lost future business, etc.), and these costs may dwarf those of the actual failure. The framework above is sufficiently flexible to account for such costs, provided they can be modeled in a realistic way that is acceptable to all parties. The framework is also sufficiently flexible to handle situations where a part is being qualified for multiple missions—and even for missions that have not yet been planned. In this case, the sum in equation (2) will include applications for all missions covered by the qualification and perhaps for likely applications of the part based on historical data from past missions.

IV. Radiation Environment Time Dependence

The time dependence of the failure cost means that it is helpful to understand the temporal variability of the applied stress (e.g. TID, DDD or LET), on which the failure distribution depends. Although the time dependence of radiation environments is still a nascent field, we can use data on past radiation-environment variability to determine a worst-case profile of stress vs. time consistent with a given confidence level. The standard SEE rate calculation methodology serves as a guide, since it yields an average SEE rate about which the observed rate fluctuates according to Poisson statistics. This allows us to bound the on-orbit rate to any desired confidence level, and the only additional time dependence results from the solar cycle. Since like the mean rate, the bound of the rate for any CL is constant in time, the failure probability is constant, and the risk will have the same time dependence as the failure cost, C_f.

The time dependence of TID or displacement damage is important mainly when a crucial part has marginal radiation performance, despite use of all practical mitigation strategies. In such cases, *when* the parts are exposed to radiation (e.g. early vs. late) may be as important for the achieving mission goals as the magnitude of radiation exposure.

We used the ESP solar proton model[8] to examine the solar-proton dominated environment for the James Webb Space Telescope (JWST). For a desired CL, we constructed a worst-case exposure vs. time profile by first generating proton fluences for that CL for missions lasting from 1 to 7 solar active years. From these, we generated incremental yearly fluences, and constructed our mission dose profile assuming the worst fluence occurs during the first year, the second worst in the second year and so on out to the 5.5 year mission duration. We fit the resulting profiles in figure 3 to a power law in time $D=D_0t^a$. As might be expected, less shielding of the application and higher required CL corresponded to greater deviation from linearity (a=1) for our worst-case profile.

For missions shorter than 2 years, individual solar particle events (SPEs) may dominate radiation exposure, so we constructed simulated missions, with SPE numbers fluctuating according to Poisson statistics and event size determined by a randomly selected CL for the ESP model. Generating 10000 such simulated missions, we determined our 95% WC mission profile as that where only 5% of profiles received more dose in the first 6 and 12 months of the mission, and subject to the constraint that the two-year dose was consistent with the two-year 95% ESP mission dose. For this method, dose $D \propto t^{1/2}$, largely independent of shielding and CL...

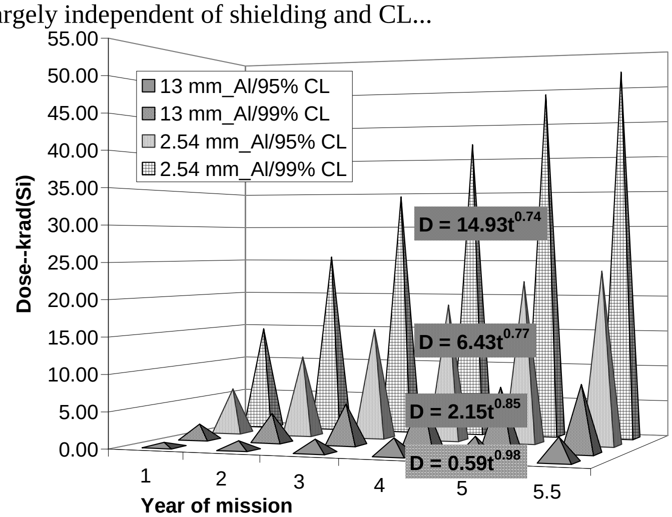

Fig. 3 TID levels at the 95% and 99% CL the ESP model for 2.54 and 13 mm (~100 and 500 mils) Al equivalent spherical shielding show that year-to-year variability is greater for higher CL and thinner equivalent shielding.

V. Prior and Posterior distributions--TID

The Prior summarizes expectations of radiation behavior before performing additional radiation testing. The preferred information source for the Prior is archival data for parts representative of the flight lot. We base our Prior on a Maximum Likelihood (ML) fit to such data. Since the logarithm of the likelihood ratio is distributed roughly as χ^2 (with degrees of freedom equal to the number of distribution fit parameters), we can construct a Prior distribution for the fit parameters from the likelihood ratio for each set of fit parameters to the ML values. For example, if μ_{BF} and σ_{BF} are the ML mean and standard deviation for a Normal fit to data, and $L(\mu_{BF},\sigma_{BF})$ is the likelihood for these parameters, then the probability for any other fit parameters, (μ,σ) will be

$$P(\mu,\sigma) = 0.5\times N\chi^2(2\ln(L(\mu,\sigma)/L(\mu_{BF},\sigma_{BF})),2) \quad (3)$$

In (3), we determine the constant N by normalization.

In Reference 1, we illustrated the use of historical data mined from the BSIS data archive to draw conclusions about RHA using a particularly large dataset (38 RLAT samples, 158 total parts) for Linear Technologies RH1014 quad op amps. The RH1014 exhibits consistent part-to-part and lot-to-lot performance. This consistency and the large amount of data available make the part a good candidate for purposes of illustration. Figure 4 shows the Prior for a Normal fit to a histogram of increased input bias current (Ibias) in RH1014s after 100 krad(Si). The good statistics produce near agreement between the best fit and 95% WC parameters.

As we add new data, we form the fit-parameter posterior distribution by updating the Prior using Bayes' theorem—

equation 1. We estimate the denominator by normalizing the Posterior distribution. If the new data are consistent with the Prior, the Posterior is little changed from the Prior. On the other hand, inconsistent data result in a Posterior significantly broader and different from the Prior, indicating that our initial hypothesis *A* is incorrect and should be revised and retested.

μ \ σ	3.70	3.85	4.00	4.15	4.30	4.45	4.60	4.75	4.90	5.05
26.35	0.00	0.01	0.02	0.05	0.07	0.07	0.05	0.02	0.01	0.00
26.50	0.00	0.02	0.08	0.15	0.19	0.17	0.11	0.06	0.02	0.01
26.65	0.01	0.07	0.20	0.36	0.43	0.36	0.23	0.11	0.04	0.01
26.80	0.03	0.14	0.41	0.71	0.81	0.66	0.39	0.18	0.07	0.02
26.95	0.05	0.25	0.68	1.15	1.28	1.00	0.58	0.27	0.10	0.03
27.10	0.07	0.35	0.93	1.52	1.66	1.28	0.74	0.33	0.12	0.04
27.25	0.08	0.38	1.02	1.66	1.80	1.38	0.79	0.35	0.13	0.04
27.40	0.07	0.34	0.90	1.48	1.62	1.26	0.72	0.32	0.12	0.04
27.55	0.05	0.24	0.65	1.09	1.22	0.96	0.56	0.26	0.09	0.03
27.70	0.03	0.13	0.38	0.66	0.76	0.62	0.37	0.17	0.07	0.02
27.85	0.01	0.06	0.18	0.32	0.39	0.33	0.21	0.10	0.04	0.01
28.00	0.00	0.02	0.07	0.13	0.17	0.15	0.10	0.05	0.02	0.01
28.15	0.00	0.01	0.02	0.04	0.06	0.06	0.04	0.02	0.01	0.00

Fig. 4 The best-fit for a Normal Prior of excess Ibias in RH1014 op amps after 100 krads has μ=27.25, σ=4.3 (shaded cells w/ diagonal stripes). The 95% WC fit is μ=28.15, σ=5.05 (shaded cross-hatched cells).

In cases where we have lots of historical data, more testing may add little precision to the probability distribution. This is often the case for SEE in space qualified parts.

VI. Prior and Posterior Distributions--SEE

The determining parameters for SEE rates are the threshold LET LET_{th}, the limiting cross section σ_{lim} and the width and shape parameters (w and s) for the Weibull fit to the σ vs. LET curve. Usually part-to-part and lot-to-lot variations in these parameters are small compared to TID variability and to the systematic errors inherent in SEE rate calculations. Thus, we are mainly interested in how observed SEE counts affect the uncertainties in LET_{th}, σ_{lim} and the Weibull fit parameters.

As for the TID case, we base our Prior on a Likelihood ratio—for a Weibull fit to historical data. However, there are subtleties in performing the ML fit. First, the fit is easier if we convert the σ vs. LET curve into a differential form—fitting differences in σ from one LET to the next to a Weibull pdf. Second, we convert each cross section into an event count by multiplying by a common fluence, selected so the sum of event counts over all effective LETs is the same as the number of events observed during the test(s) on which the dataset is based. This ensures the likelihood fit has the appropriate sensitivity to additional data. Figure 5 shows the Prior and the resulting best, 95% WC and 99% WC Weibull fits to a SEL σ vs. LET curve for the XA1 ASIC, used in the Burst-Alert Telescope (BAT) of NASA's SWIFT gamma-ray burst observatory.[9] Because the BAT uses 128 XA1s to process signals from SWIFT's gamma-ray detectors, over 500 SELs were recorded in an effort to minimize statistical errors and to find rare destructive SEL modes. This large dataset produces a Prior with small error contours on the fit parameters.

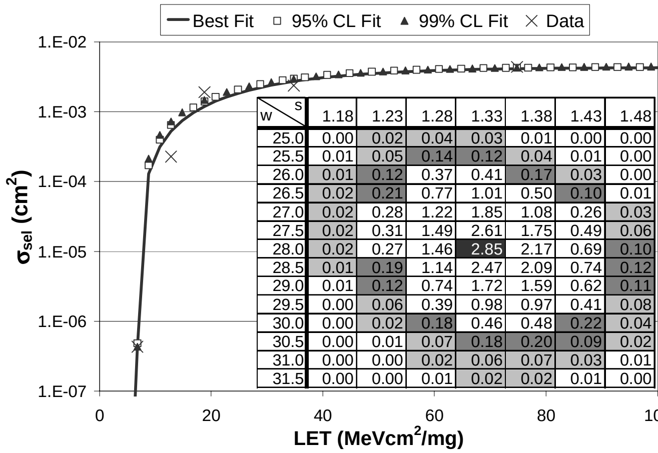

w \ s	1.18	1.23	1.28	1.33	1.38	1.43	1.48
25.0	0.00	0.02	0.04	0.03	0.01	0.00	0.00
25.5	0.01	0.05	0.14	0.12	0.04	0.01	0.00
26.0	0.01	0.12	0.37	0.41	0.17	0.03	0.00
26.5	0.02	0.21	0.77	1.01	0.50	0.10	0.01
27.0	0.02	0.28	1.22	1.85	1.08	0.26	0.03
27.5	0.02	0.31	1.49	2.61	1.75	0.49	0.06
28.0	0.02	0.27	1.46	2.85	2.17	0.69	0.10
28.5	0.01	0.19	1.14	2.47	2.09	0.74	0.12
29.0	0.01	0.12	0.74	1.72	1.59	0.62	0.11
29.5	0.00	0.06	0.39	0.98	0.97	0.41	0.08
30.0	0.00	0.02	0.18	0.46	0.48	0.22	0.04
30.5	0.00	0.01	0.07	0.18	0.20	0.09	0.02
31.0	0.00	0.00	0.02	0.06	0.07	0.03	0.01
31.5	0.00	0.00	0.01	0.02	0.02	0.01	0.00

Fig. 5 The inset shows the Prior for a Weibull fit to a SEL σ vs. LET curve for the XA1 ASIC, along with the best-fit parameters (black cell), and contours for the 95% CL (dark gray) and 99% CL (light gray). In the main figure, due to the large dataset, the best-fit (solid line) to the data (x's), and the 95% WC (open boxes) and 99% WC (solid triangles) fits nearly coincide.

If SEE susceptibility is consistent from part to part and lot to lot (usually the case for space and military parts), additional testing adds little precision to the calculated rate. (In SWIFT's 600 km, 22° inclination orbit, the SEL rate would be 1.26×10^{-4} SELs per device-day for the best fit, with the 95% WC and 99% WC rates being 17% and 23% higher, respectively). In such cases parts are commonly qualified based solely on historical data. A Bayesian analysis provides a justification for this approach, because the added risk reduction would not justify the test cost. If testing is done, we use the Prior to determine how much more testing is needed (see below) and then update the Prior with the new data using Bayes' Theorem.

VII. Applications: TID Testing Decisions

Many programs seek to save money by waiving RLAT for parts with a long and consistent TID performance history. While this strategy poses risks (e.g., due to process changes or "bad lots"), it can redirect resources where they will most benefit the project. For parts such as the RH1014, manufacturer's guarantees and consistent performance history suggest that the risks could be addressed by periodic testing, provided designs include sufficient margin.

In reference 1, we proposed a method for estimating design-to values for critical parameters based on application needs and a given acceptable failure probability. The analogous Bayesian treatment emerges somewhat more naturally in that an acceptable failure probability need not be assigned, but emerges from the requirement that we test only if the resultant risk reduction exceeds the test cost C_t. If $R_f<C_t$, additional testing will not reduce risk. For constant C_f, R_f will be less than C_t if the cumulative failure probability $P_f<C_t/C_f$. For a hypothetical application assuming a 99% CL, 7 year ESP environment transported through 1.3 mm Al equivalent shielding, the mission TID is 215 krad(Si) and the 99% worst-case dose vs. time profile follows a power law with $D\propto t^{2/3}$. If C_f=\$1000000 and C_t=\$10000, then we meet our criterion for waiving RLAT when $P_f\leq1\%$ integrated over the mission

duration (see figure 6). The design-to leakage current, I_d, is the value for which P(Ibias>I_d)<1%--in this case 95 nA. For the RH1014 and many other parts, I_d and μ increase roughly linearly with dose from 0-200 krad(Si). Multiplying the ratio I_d/μ=1.66 by the usual 2x radiation design margin (RDM), the RDM for 95/99 assurance is roughly 3.3×.

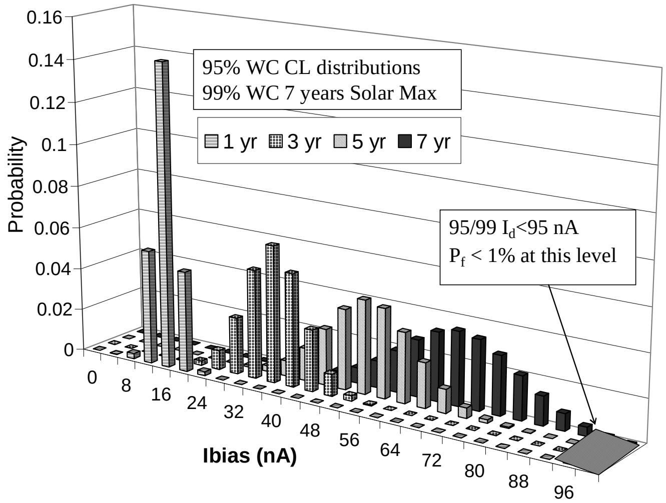

Fig. 6 The 95% CL worst-case Ibias pdf for RH 1014s exposed to a 99% CL ESP environment behind 1.3 mm Al shifts upward and broadens during 7 years of Solar Max. Setting the design-to Ibias (Id) to 95 nA, P(Ibias>I_d) (shaded rectangle) is less than 1%. This defines the 95/99 CL I_d.

If C_f is time dependent (e.g. decreasing linearly with time) the calculation is more complicated, but the criterion (R_f<C_t) is the same. Again for $C_f(t$=0)=$1000000 and C_t=$10000, figure 7 illustrates that introducing the time dependence of C_f reduces I_d to 84 nA and the required RDM to 2.95x.

Regardless of the time dependence of C_f, the large dataset and reasonable radiation response of the RH1014 allow us to substitute moderate RDMs for RLAT. While no historical data can protect against adverse process changes, this approach can at least justify substituting periodic testing for RLAT in some cases. Smaller datasets require larger RDMs.

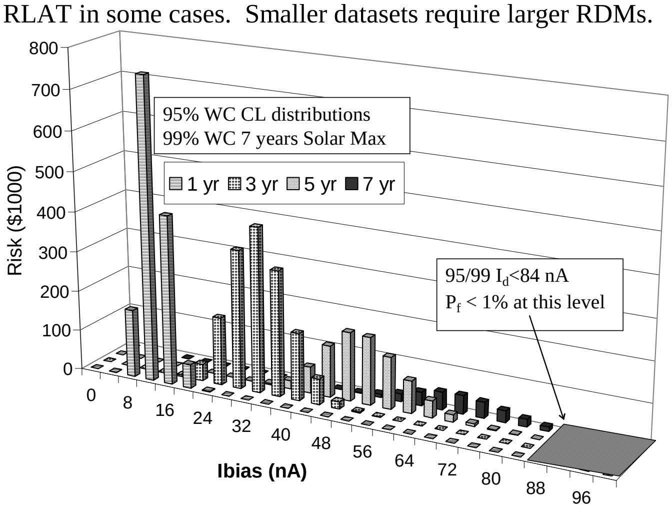

Fig. 7 At the design-to Ibias, I_d, integrated failure risk R_f equals the test cost C_t, so more testing does not reduce risk. A failure cost that decreases linearly with time lowers I_d by about 12% compared to a constant failure cost.

VIII. Applications: SEE Testing Decisions

As alluded to above, if we have high confidence that flight parts will behave like those on which the Prior is based, it may be possible to qualify a part based on the Prior alone. Since the SEE probability density is constant, any time dependence in equation 2 arises from the failure cost, C_f. For constant C_f, the risk is C_f multiplied by the expected number of SEE during the mission. For linearly decreasing C_f, we would multiply the mission SEE count by the average C_f (half the initial value). For the example of the XA1 ASICs discussed above, we expect fewer than 12 SELs (95% CL) during SWIFT's 2 year mission. Since the mission could suffer 10% attrition and still fulfill its requirements, and since fewer than 3% of the SELs during testing resulted in device failure even at the highest test LET, we concluded that the SEL susceptibility does not pose a serious threat to mission requirements. The consequences of nondestructive SELs were mitigated with SEL detection and recovery circuitry.

Now, we investigate the use of the Prior to determine whether additional testing is warranted. Let us suppose that instead of the 500 SELs in our XA1 dataset, our Prior is instead based on 50 observed SELs, plus 2 destructive SELs observed in a special test at high LET (see Table I).

Table I: SEL Dataset for XA1 SEL Prior

LET ($MeVcm^2/mg$)	# SEL (test current limited)	σ_{sel} Total (cm^2)	Destructive SEL	σ_{dest} (cm^2)
6.8	2	4.3E-07	N/A	N/A
11.4	8	0.000228	N/A	N/A
18.7	12	0.0019	N/A	N/A
26.2	N/A	N/A	0	<0.0000002
37.4	13	0.00237	N/A	N/A
59.9	N/A	N/A	2	0.0000012
74.8	15	0.00438	N/A	N/A

The lower amount of data results in larger error bars on the fit parameters and the resulting rates. (See figure 8.) Also, assume the XA1 is being considered for a two-year mission in geostationary orbit (GEO)—for which SEL rates are roughly 45 times those for the SWIFT orbit. We wish to address both destructive SEL and the impact of nondestructive SELs.

With regard to destructive events, we want to ensure that we have 128 ASICs operational at end of life. While the ASICs will have overcurrent detection and recovery circuitry, we cannot be certain that this circuitry will protect them against all destructive events. We therefore seek to meet our goals by adding redundant ASICs. Specifically, we want to know how many cold spare XA1s we must fly to have 95% confidence of at least 99% probability that at least 128 ASICs will be operational at the end of the 2 year mission. Because the finite supply of ASICs precludes gathering more destructive SEL data, we calculate a bounding rate using a figure of merit (FOM) approach—$R \approx 400\sigma_{dest}/LET_{th}^2$. Taking the 95% CL bounding cross section (assuming Poisson fluctuations)—3.7×10^{-6} cm^2, and LET_{th}= 26 $MeVcm^2/mg$ yields a 95% CL rate of 2.3×10^{-6} SELs per device-day, or an expected mission total of 0.2.destructive SELs. Since a Poisson distribution with this mean has less than 1% probability of more than 2 SELs, 2 cold spares provide the required level of assurance.

w \ s	0.93	1.13	1.33	1.53	1.73	1.93
18.0	0.00	0.00	0.00	0.00	0.00	0.00
20.0	0.00	0.01	0.00	0.00	0.00	0.00
22.0	0.01	0.04	0.03	0.00	0.00	0.00
24.0	0.01	0.09	0.12	0.03	0.00	0.00
26.0	0.01	0.13	0.24	0.09	0.01	0.00
28.0	0.01	0.13	0.29	0.14	0.02	0.00
30.0	0.00	0.09	0.24	0.14	0.03	0.00
32.0	0.00	0.06	0.15	0.09	0.02	0.00
34.0	0.00	0.03	0.07	0.04	0.01	0.00
36.0	0.00	0.01	0.03	0.02	0.00	0.00
38.0	0.00	0.01	0.01	0.01	0.00	0.00
40.0	0.00	0.00	0.00	0.00	0.00	0.00
42.0	0.00	0.00	0.00	0.00	0.00	0.00

Fig. 8 Diminishing the dataset on which the Prior in figure 5 is based by 10x (for a constant σ vs. LET curve constant, but decreased fluences) broadens the 95% and 99% confidence contours (dark and light grey, respectively), while the best-fit parameters (black square) stay the same.

The possibility of destructive SEL complicates the recovery process for nondestructive events, since the affected device must be verified functional before full operations are recovered. The loss of availability during this time incurs a cost of \$300. Using the Prior (figure 8), the expected mission total of SELs for the best-fit Weibull (w=28, s=1.33) is 531, while the 95% WC fit (w=21, s=1) expected total is 832. Thus, the 95% WC risk is \$249600. This is clearly larger than the likely cost of a subsequent SEL test, leading to the question of whether such a test is likely to reduce risk sufficiently to justify its cost. As mentioned above, if the test results support the Prior, the posterior will be narrower than the Prior and have roughly the same best-fit. In this case, the 95% WC risk will approach the best-fit value as data are added. To simulate this process, we repeatedly add data consistent with the dataset underlying the Prior and observe the process of convergence. However, since data is not a free commodity, we assume that every SEL added to the dataset costs an average of \$100 (in test preparation, beam time, etc.), on top of a minimum test cost of \$2000. While more sophisticated analyses are possible, this procedure is adequate for test planning. With these assumptions, we determine the optimum test in terms of data added. As figure 11 shows, increasing the dataset to 229 SELs over the 50 for the Prior at a cost of \$19900 reduces risk by \$56779, for a net risk reduction of \$36879. While the data accumulated would likely vary somewhat from the data underlying the Prior, the symmetry of the Prior indicates that the chances the new data will be worse than expected are roughly equal to the chances that it will be better. Thus, our analysis establishes a good probability that additional testing would reduce risk and that a level of effort approaching \$20000 is appropriate. Moreover, the analysis can be updated in real time during testing, and the testing strategy can be optimized accordingly.

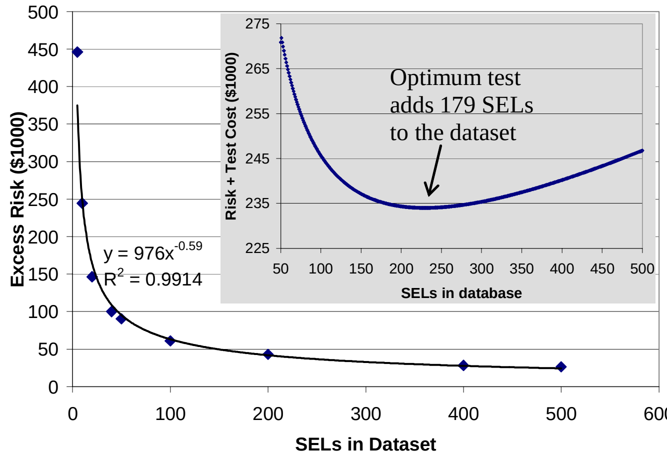

Fig. 11 As we add data, the excess 95% WC risk over that for the best-fit to the data approaches 0—roughly according to a power law (main figure). Since testing is costly, there is an optimum test effort (inset) that minimizes the combined risk + test cost—here adding 179 SELs at a cost of \$19900.

IX. Priors with Limited Data

Bayesian methods allow us to develop Priors using expert opinion, modeling, etc. to supplement, or even replace, data. We validate our assumptions by updating the Prior with test data. Regardless of its underlying source, the Prior should reflect our model's uncertainty, and tests should distinguish between our model and any alternatives (A and ~A in Bayes' Theorem). If $P(B|A) > P(B|{\sim}A)$, the posterior $P(A|B)$ will more clearly favor hypothesis A. If B conflicts with A, the posterior becomes more diffuse, suggesting our model is incorrect and needs more testing.

When data are scarce, the Prior is useful for test planning and resource allocation. To illustrate, we consider the TID response of CD4000 family parts baselined for control units for JWST optics. These parts must operate at 25 Kelvin for 5.5 years in the L2 radiation environment (22 krad(Si) @ 95% CL by EOL). Although CD4000s perform well to >100 krad(Si) at ambient temperatures, similar performance at 25 K is questionable, given the suppression of interface state formation and annealing at cryogenic temperatures.

In testing to date, the CD4051 mux and CD4066 switch revealed considerable variability, with all 5 CD4051s failing parametrically between 10 and 20 krad(Si), one CD4066 failing at 30 krad(Si), another at 70 krad(Si) and the other 3 still functional when testing was suspended at 75 krad(Si).

These results raised concerns for the hardness of the CD40109 level shifter—still to be tested. Failure to qualify this part would result in redesign with a significant mass penalty. We define failure as leakage current exceeding specifications and assume C_f decreases linearly in time: C_f=\$1E7×(1-t/5.5). C_t for a cryogenic TID test is on the order of \$15000-\$20000. Constructing a Prior for the CD4066 from the previous cryogenic TID data is straightforward, although the broad Prior makes a Normal approximation unphysical.

Table II shows TID as a function of time (1st 2 columns--95% WC dose vs. time profile), the failure cost, P_f for the Prior (3rd column--90% CL) and P_f with 5 and 8 more parts failing above 75 krad(Si) (last 2 columns—90% CL). Other guidelines from studying the Prior and its response to new data sets are:

1) Limiting the failure distribution breadth (given the 30 krad(Si) failure) requires testing the next lot to failure.

2) 5 parts are adequate to attain 90/99 statistics for a Weibull distribution; 8 are needed if the distribution looks lognormal. R_f for the cost and failure probability profiles in columns 3 and 5 is $9800, so more testing above the 5-8 additional parts would not be justified.

TABLE II: CD4066 90% CL CUMULATIVE WEIBULL P_F

Mission Year	TID krad(Si)	Failure Cost	P(Fail)-Prior	P(Fail)-Prior + 5 parts >75 krad(Si)	P(Fail)-Prior + 8 parts >75 krad(Si)
1	6.14	8.18E+06	0.016	0.00007	0
2	10.7	5.21E+06	0.095	0.0005	0.00001
3	14.6	2.37E+06	0.125	0.0017	0.00005
4	17.7	6.45E+05	0.145	0.0035	0.00015
5	20.7	5.87E+04	0.166	0.006	0.0004
5.5	22	0.00E+00	0.175	0.008	0.0005

The different results for the CD4066 and CD4051 complicate generation of a Prior for the CD40109. Based on cryogenic performance data and expert opinion, we construct the Prior weighting results for the CD4066 by 90% and the CD4051 by 10%. We give CD4000 family parts half the weight for CD40109 data. The resulting Prior in figure 10 shows that broad failure distributions with high mean hardness are as probable as narrow distributions peaked 70-90 krad(Si). Again, testing to failure is important to narrow the distribution.

Table III gives the CD40109 P_f in a format similar Table I. Adding data (5 parts w/ failure >75 krad(Si)), narrows the distribution in both Weibull width, W, and shape, S. (See Fig. 11.) Again, 5 parts with failure levels >75 krad(Si) are adequate to establish 90/99 statistics if the failure distribution resembles a Weibull, while 8 parts are needed if the distribution looks more lognormal. For this distribution, R_f =13000, again obviating the need for further testing if the P_f in column 5 is achieved or exceeded.

W \ S	0.20	0.95	1.70	2.45	3.20	3.95	4.70	5.45	6.20	6.95
25	0.02	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
37	0.03	0.05	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00
49	0.03	0.10	0.06	0.02	0.00	0.00	0.00	0.00	0.00	0.00
61	0.03	0.15	0.15	0.10	0.06	0.03	0.01	0.00	0.00	0.00
73	0.03	0.18	0.22	0.19	0.13	0.09	0.06	0.03	0.02	0.01
85	0.04	0.20	0.26	0.22	0.16	0.11	0.07	0.04	0.02	0.01
97	0.04	0.22	0.27	0.22	0.15	0.09	0.05	0.02	0.01	0.01
109	0.04	0.23	0.26	0.20	0.12	0.06	0.03	0.01	0.01	0.00
121	0.04	0.23	0.25	0.17	0.09	0.04	0.02	0.01	0.00	0.00
133	0.04	0.23	0.23	0.14	0.07	0.03	0.01	0.00	0.00	0.00
145	0.04	0.23	0.22	0.12	0.05	0.02	0.01	0.00	0.00	0.00
157	0.04	0.23	0.20	0.10	0.04	0.01	0.00	0.00	0.00	0.00
169	0.04	0.23	0.18	0.08	0.03	0.01	0.00	0.00	0.00	0.00
181	0.04	0.22	0.16	0.07	0.02	0.01	0.00	0.00	0.00	0.00

Fig. 10 Generic CD4000 family Prior Weibull distribution, weighted 90% on theCD4066 and 10% on the CD4051.

TABLE III: CD40109 90% CL CUMULATIVE WEIBULL P_F

Mission Year	TID krad(Si)	Failure Cost	P(Fail)-Prior	P(Fail)-Prior + 5 parts >75 krad(Si)	P(Fail)-Prior + 8 parts >75 krad(Si)
1	6.14	8.18E+06	0.115	0.0001	0
2	10.7	5.21E+06	0.175	0.0007	0.000007
3	14.6	2.37E+06	0.22	0.0022	0.00005
4	17.7	6.45E+05	0.25	0.0045	0.00015
5	20.7	5.87E+04	0.275	0.008	0.0004
5.5	22	0.00E+00	0.28	0.01	0.0006

W \ S	2.45	3.20	3.95	4.70	5.45	6.20	6.95	7.70	8.45	9.20	9.95	10.70
65	0.01	0.02	0.02	0.02	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00
67.5	0.01	0.03	0.05	0.05	0.05	0.03	0.02	0.01	0.01	0.00	0.00	0.00
70	0.02	0.04	0.08	0.10	0.11	0.11	0.09	0.07	0.05	0.04	0.02	0.01
72.5	0.02	0.05	0.10	0.15	0.18	0.20	0.20	0.19	0.16	0.13	0.10	0.08
75	0.02	0.06	0.11	0.17	0.22	0.25	0.26	0.25	0.23	0.20	0.17	0.13
77.5	0.02	0.06	0.11	0.17	0.21	0.24	0.24	0.22	0.19	0.16	0.13	0.10
80	0.02	0.06	0.10	0.15	0.18	0.18	0.17	0.14	0.11	0.09	0.06	0.04
82.5	0.02	0.05	0.09	0.12	0.13	0.12	0.10	0.08	0.05	0.04	0.02	0.01
85	0.02	0.05	0.07	0.08	0.08	0.07	0.05	0.03	0.02	0.01	0.01	0.00
87.5	0.02	0.04	0.05	0.06	0.05	0.04	0.02	0.01	0.01	0.00	0.00	0.00
90	0.02	0.03	0.04	0.04	0.03	0.02	0.01	0.01	0.00	0.00	0.00	0.00
92.5	0.01	0.02	0.03	0.02	0.02	0.01	0.00	0.00	0.00	0.00	0.00	0.00
95	0.01	0.02	0.02	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00
97.5	0.01	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Fig 11 CD40109 Posterior Weibull distribution, assuming the Prior from figure 10 and 5 parts failing at levels >75 krad(Si).

X. CONCLUSION

The above analyses have provided a cursory introduction to the broad applicability of Bayesian analysis to RHA issues—both for SEE and degradation phenomena such as TID (as well as displacement damage). Bayesian risk serves as both a metric for allocating resources, judging RHA efficacy and comparing risks and as a rigorous and efficient way of using diverse information. However, Bayesian methods while rigorous, do not banish subjectivity. Subjectivity arises both in defining the Prior—e.g. the information to include, the weights for different data, the models assumed—and in defining failure costs. In our analyses, we have taken a conservative approach to both of these aspects of the analysis. Priors were defined in terms of archival data representative of the flight lot—or in the case of the CD40109, in terms of closely related parts from the same radiation hardened logic family. Failure costs were defined in terms of either contractual incentives or penalties or in terms of functions compromised by the error, degradation or failure. Since all subsequent probabilities are conditional on the Prior, and since the failure, test and remediation costs determine the scale of the RHA effort, conservative but realistic treatments of these issues are more likely to lead to conclusions that are bounding but not punitive. By making subjectivity explicit in the analysis, we can validate it and correct it as needed—an iterative process that reflects the ongoing nature of RHA.

We have also introduced significant flexibility in terms of a method for bounding at a given CL the temporal variations in the radiation environment. This allows us to introduce time-dependent failure costs and risks. This strategy allows us to achieve greater fidelity to the actual risks faced by the mission. In addition to the examples outlined above, we can foresee several other applications of Bayesian methods. The possibility in Bayesian analysis to construct a Prior based on information other than test data provides opportunities to introduce Priors based on modeling and simulations. One particular area of interest is the analysis of relatively rare, but high impact failure/error modes such as severe ELDRS response and extremely long transients in analog devices.[10]

We also note that Bayesian methods should also prove fruitful in analysis of other spacecraft threats—from reliability to mechanical failure to spacecraft charging. This prospect raises the possibility that the risks posed by these threats might be evaluated and compared—as long as all of the analyses were conducted consistent assumptions and comparable conservatism.

It is also appropriate to consider what the methods outlined here cannot do. The sources of information underlying the Prior—historical data, expert opinion etc.—are unlikely to account for unexpected threats (e.g. process changes, counterfeit parts, and so on). Also, the flexibility and power of Bayesian methods entails some risks—whether they arise from deliberate abuse or merely from inconsistent assumptions within or between analyses. Bayesian methods are most likely to find application where there is either significant amounts of information other than lot-specific test data or where test data by itself is unlikely to answer all questions about a part's suitability for its application. The techniques outlined here are complementary to conventional RHA methods,[11] rather than a replacement for them.

Finally, we have taken a Bayesian approach in the current work because we feel that such an approach provides a more natural framework for the techniques presented here. However, we note that much of the work we have discussed can also be interpreted in terms of the more common frequentist approach to statistical analysis.

REFERENCES

[1] R. Ladbury and J. Gorelick, Nuclear and Space Radiation Effects Conference July 11-15, 2005, Seattle, WA

[2] E. S. Yudkowsky, "An Intuitive Explanation of Bayesian Reasoning," http://yudkowsky.net/bayes/bayes.html . September 1, 2005.

[3] A. Namenson, "Lot Uniformity and Small Sample Sizes in Hardness Assurance", IEEE Trans. Nucl. Sci. 35, No. 6, Dec. 1988.

[4] J. von Plato, *A History of Modern Probability*, Cambridge University Press, Cambridge, UK, 1998.

[5] M. Ramoni and P. Sebastiani, "Bayesian Methods,".in *Intelligent Data Analysis*, M. Berthold and D. Hand, eds., Springer-Verlag: Berlin, 2003.

[6] P. H. Garthwaite, I. T. Jolliffe and B. Jones, *Statistical Inference*, Prentice Hall London, 1995.

[7] K. LaBel et al., "Single Event Criticality Analysis", available at http://radhome.gsfc.nasa.gov/radhome/papers/seecai.htm.

[8] M. A. Xapsos, "Model for Solar Proton Risk Assessment," IEEE Trans. Nucl Sci., 51, No. 6, Dec. 2004.

[9] M. O'bryan et al., "Current Single Event Effects and Radiation Damage Results for Candidate Spacecraft Electronics," 2002 IEEE Radiation Effects Data Workshop, 15-19 July 2002.

[10] Y. Boulghassoul, et al., "Investigation of Millisecond-Long Analog Single-Event Transients in the LM6144 Op Amp", IEEE Trans. Nucl Sci., 51, No. 6, Dec. 2004.

[11] R. Pease, "Microelectronic Piece Part Radiation Hardness Assurance for Space Systems", 2004 NSREC Short Course.